

EAST Search History

Ref #	Hits	Search Query	DBs	Default Operator	Plurals	Time Stamp
L1	27524	"711"/\$.ccls.	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2006/01/07 14:11
L2	11	(Jan near2 gray).in.	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2006/01/07 14:11
L3	3780	quer\$6 same cache	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2006/01/07 14:11
L5	119176	semiconductor near2 memory	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2006/01/07 14:13
L6	887	(spatial same cache)	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2006/01/07 14:14
L7	1913	(locality same cache)	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2006/01/07 14:14
L8	295	multilevel adj cache	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2006/01/07 14:15
L9	363	hierarchical adj cache	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2006/01/07 14:15
L10	37796	access adj time	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2006/01/07 14:15

EAST Search History

L11	684369	(determin\$5 or estimat\$4) with time	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2006/01/07 14:17
L12	104	3 and 5	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2006/01/07 14:17
L13	573	6 and 7	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2006/01/07 14:17
L14	0	12 and 13	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2006/01/07 14:17
L15	640	8 or 9	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2006/01/07 14:17
L16	0	15 and 12	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2006/01/07 14:18
L17	25	15 and 13	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2006/01/07 14:18
L18	710765	10 or 11	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2006/01/07 14:18
L19	15	17 and 18	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2006/01/07 14:18



Welcome United States Patent and Trademark Office

[Search Session History](#)

[BROWSE](#)

[SEARCH](#)

[IEEE XPLORE GUIDE](#)

Sat, 7 Jan 2006, 2:33:16 PM EST

Edit an existing query or compose a new query in the Search Query Display.

Search Query Display

Select a search number (#) to:

- Add a query to the Search Query Display
- Combine search queries using AND, OR, or NOT
- Delete a search
- Run a search

Recent Search Queries

- #1

((semiconductor memory)<in>metadata)
- #2

(hiearchical cache<IN>metadata)
- #3

(multilevel cache<in>metadata)
- #4

(open cache access<IN>metadata)
- #5

(spatial and locality<IN>metadata)
- #6

((hiearchical cache<IN>metadata)) and <AND> ((multilevel cache<in>metadata))
- #7

((hiearchical cache<IN>metadata)) <AND> ((spatial and locality<IN>metadata))
- #8

access adj latency
- #9

(access latency<IN>metadata)
- #10

((spatial and locality<IN>metadata)) <AND> ((access latency<IN>metadata))




[Subscribe \(Full Service\)](#) [Register \(Limited Service, Free\)](#) [Login](#)

 Search: ☒ The ACM Digital Library ☐ The Guide


[Feedback](#) [Report a problem](#) [Satisfaction survey](#)

Terms used

[spatial](#) [locality](#) [cache](#) [memory](#) [hierarchical](#) [multilevel](#)

Found 60 of 169,166

Sort results by


[Save results to a Binder](#)
[Try an Advanced Search](#)
[Try this search in The ACM Guide](#)

Display results


[Search Tips](#)
☐ Open results in a new window

Results 1 - 20 of 60

Result page: [1](#) [2](#) [3](#) [4](#) [next](#)Relevance scale ☐ ☐ ☐ ☐ ☐

1 [External memory algorithms and data structures: dealing with massive data](#)



Jeffrey Scott Vitter

June 2001 **ACM Computing Surveys (CSUR)**, Volume 33 Issue 2

Publisher: ACM Press

Full text available: [pdf\(828.46 KB\)](#)
 Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

Data sets in large applications are often too massive to fit completely inside the computers internal memory. The resulting input/output communication (or I/O) between fast internal memory and slower external memory (such as disks) can be a major performance bottleneck. In this article we survey the state of the art in the design and analysis of external memory (or EM) algorithms and data structures, where the goal is to exploit locality in order to reduce the I/O costs. We consider a varie ...

Keywords: B-tree, I/O, batched, block, disk, dynamic, extendible hashing, external memory, hierarchical memory, multidimensional access methods, multilevel memory, online, out-of-core, secondary storage, sorting

2 [Implications of hierarchical N-body methods for multiprocessor architectures](#)



Jaswinder Pal Singh, John L. Hennessy, Anoop Gupta

May 1995 **ACM Transactions on Computer Systems (TOCS)**, Volume 13 Issue 2

Publisher: ACM Press

Full text available: [pdf\(4.66 MB\)](#)
 Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#), [review](#)

To design effective large-scale multiprocessors, designers need to understand the characteristics of the applications that will use the machines. Application characteristics of particular interest include the amount of communication relative to computation, the structure of the communication, and the local cache and memory requirements, as well as how these characteristics scale with larger problems and machines. One important class of applications is based on hierarchical N-body methods, w ...

Keywords: N-body methods, communication abstractions, locality, message passing, parallel applications, parallel computer architecture, scaling, shared address space, shared memory

3 A locality-preserving cache-oblivious dynamic dictionary

Michael A. Bender, Ziyang Duan, John Iacono, Jing Wu


January 2002 **Proceedings of the thirteenth annual ACM-SIAM symposium on Discrete algorithms**

Publisher: Society for Industrial and Applied Mathematics

Full text available:  [pdf\(1.06 MB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#)

This paper presents a simple dictionary structure designed for a hierarchical memory. The proposed data structure is *cache oblivious* and *locality preserving*. A cache-oblivious data structure has memory performance optimized for all levels of the memory hierarchy even though it has no memory-hierarchy-specific parameterization. A locality-preserving dictionary maintains elements of similar key values stored close together for fast access to ranges of data with consecutive keys. The d ...

4 Locality optimizations for multi-level caches

 Gabriel Rivera, Chau-Wen Tseng

January 1999 **Proceedings of the 1999 ACM/IEEE conference on Supercomputing (CDROM)**

Publisher: ACM Press

Full text available:  [pdf\(176.48 KB\)](#) Additional Information: [full citation](#), [references](#), [citations](#), [index terms](#)

5 An educational tool for testing hierarchical multilevel caches

 J. A. Gómez Pulido, J. M. Sánchez Pérez, J. A. Moreno Zamora

September 1996 **ACM SIGARCH Computer Architecture News**, Volume 24 Issue 4


Publisher: ACM Press

Full text available:  [pdf\(574.53 KB\)](#) Additional Information: [full citation](#), [abstract](#), [index terms](#)

In this work we present a simulator for a multilevel cache memory system on a monoprocessor environment. It has incorporated a full graphic interface operating on a PC-DOS environment. At first, the simulator was conceived as a tool for applying it to teaching of cache memories. However, the potentiality of the developed system has proved its utility on program analysis and design strategies of memory systems. The above characteristics enable the simulator to be used for designing systems that r ...

Keywords: education, multilevel caches, performance evaluation, trace-driven simulation

6 Improving data locality with loop transformations

 Kathryn S. McKinley, Steve Carr, Chau-Wen Tseng

July 1996 **ACM Transactions on Programming Languages and Systems (TOPLAS)**, Volume 18 Issue 4

Publisher: ACM Press

Full text available:  [pdf\(411.40 KB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#), [review](#)

In the past decade, processor speed has become significantly faster than memory speed. Small, fast cache memories are designed to overcome this discrepancy, but they are only effective when programs exhibit data locality. In the this article, we present compiler optimizations to improve data locality based on a simple yet accurate cost model. The model computes both temporal and spatial reuse of cache lines to find desirable loop organizati ...

Keywords: Cache, compiler optimization, data locality, loop distribution, loop fusion, loop permutation, loop reversal, loop transformations, microprocessors, simulation

7 Large models & large displays: Cache-oblivious mesh layouts

Sung-Eui Yoon, Peter Lindstrom, Valerio Pascucci, Dinesh Manocha
July 2005 **ACM Transactions on Graphics (TOG)**, Volume 24 Issue 3

Publisher: ACM Press

Full text available: [pdf\(447.02 KB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

We present a novel method for computing cache-oblivious layouts of large meshes that improve the performance of interactive visualization and geometric processing algorithms. Given that the mesh is accessed in a reasonably coherent manner, we assume no particular data access patterns or cache parameters of the memory hierarchy involved in the computation. Furthermore, our formulation extends directly to computing layouts of multi-resolution and bounding volume hierarchies of large meshes. We deve ...

8 Automatic tiling of iterative stencil loops

Zhiyuan Li, Yonghong Song
November 2004 **ACM Transactions on Programming Languages and Systems (TOPLAS)**, Volume 26 Issue 6

Publisher: ACM Press

Full text available: [pdf\(947.69 KB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

Iterative stencil loops are used in scientific programs to implement relaxation methods for numerical simulation and signal processing. Such loops iteratively modify the same array elements over different time steps, which presents opportunities for the compiler to improve the temporal data locality through loop tiling. This article presents a compiler framework for automatic tiling of iterative stencil loops, with the objective of improving the cache performance. The article first presents a ...

Keywords: Caches, loop transformations, optimizing compilers

9 Compiler-based I/O prefetching for out-of-core applications

Angela Demke Brown, Todd C. Mowry, Orran Krieger
May 2001 **ACM Transactions on Computer Systems (TOCS)**, Volume 19 Issue 2

Publisher: ACM Press

Full text available: [pdf\(499.03 KB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#), [review](#)

Current operating systems offer poor performance when a numeric application's working set does not fit in main memory. As a result, programmers who wish to solve "out-of-core" problems efficiently are typically faced with the onerous task of rewriting an application to use explicit I/O operations (e.g., read/write). In this paper, we propose and evaluate a fully automatic technique which liberates the programmer from this task, provides high performance, and requires only minima ...

Keywords: compiler optimization, prefetching, virtual memory

10 Register tiling in nonrectangular iteration spaces

Marta Jiménez, José M. Llabería, Agustín Fernández
July 2002 **ACM Transactions on Programming Languages and Systems (TOPLAS)**, Volume 24 Issue 4

Publisher: ACM Press

Full text available: [pdf\(2.21 MB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

Loop tiling is a well-known loop transformation generally used to expose coarse-grain parallelism and to exploit data reuse at the cache level. Tiling can also be used to exploit

data reuse at the register level and to improve a program's ILP. However, previous proposals in the literature (as well as commercial compilers) are only able to perform multidimensional tiling for the register level when the iteration space is rectangular. In this article we present a new general algorithm to perform m ...

Keywords: Data reuse, locality, loop optimization, loop tiling, register level

11 Competitive algorithms for multilevel caching and relaxed list update

Marek Chrobak, John Noga

January 1998 **Proceedings of the ninth annual ACM-SIAM symposium on Discrete algorithms**

Publisher: Society for Industrial and Applied Mathematics

Full text available:  pdf(1.10 MB) Additional Information: [full citation](#), [references](#), [citations](#), [index terms](#)



12 System-level power optimization: techniques and tools



Luca Benini, Giovanni de Micheli

April 2000 **ACM Transactions on Design Automation of Electronic Systems (TODAES)**, Volume 5 Issue 2

Publisher: ACM Press

Full text available:  pdf(385.22 KB) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)



This tutorial surveys design methods for energy-efficient system-level design. We consider electronic systems consisting of a hardware platform and software layers. We consider the three major constituents of hardware that consume energy, namely computation, communication, and storage units, and we review methods of reducing their energy consumption. We also study models for analyzing the energy cost of software, and methods for energy-efficient software design and compilation. This survey ...

13 Energy-aware design of embedded memories: A survey of technologies, architectures, and optimization techniques



Luca Benini, Alberto Macii, Massimo Poncino

February 2003 **ACM Transactions on Embedded Computing Systems (TECS)**, Volume 2 Issue 1

Publisher: ACM Press

Full text available:  pdf(288.44 KB) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)



Embedded systems are often designed under stringent energy consumption budgets, to limit heat generation and battery size. Since memory systems consume a significant amount of energy to store and to forward data, it is then imperative to balance power consumption and performance in memory system design. Contemporary system design focuses on the trade-off between performance and energy consumption in processing and storage units, as well as in their interconnections. Although memory design is as ...

Keywords: Embedded systems, embedded memories, integration, memories, nonvolatile, system-on-a-chip, volatile

14 An accurate cost model for guiding data locality transformations



Xavier Vera, Jaume Abella, Josep Llosa, Antonio González

September 2005 **ACM Transactions on Programming Languages and Systems (TOPLAS)**, Volume 27 Issue 5

Publisher: ACM Press

Full text available:  pdf(2.50 MB) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)



Caches have become increasingly important with the widening gap between main memory and processor speeds. Small and fast cache memories are designed to bridge this discrepancy. However, they are only effective when programs exhibit sufficient data locality. The performance of the memory hierarchy can be improved by means of data and loop transformations. Tiling is a loop transformation that aims at reducing capacity misses by shortening the reuse distance. Padding is a data layout transformation ...

Keywords: Cache memories, genetic algorithms, padding, tiling

15 A comparison of three programming models for adaptive applications on the Origin2000

Hongzhang Shan, Jaswinder P. Singh, Leonid Oliker, Rupak Biswas

November 2000 **Proceedings of the 2000 ACM/IEEE conference on Supercomputing (CDROM)**

Publisher: IEEE Computer Society

Full text available:  [pdf\(239.30 KB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)
 [Publisher Site](#)

Adaptive applications have computational workloads and communication patterns which change unpredictably at runtime, requiring load balancing to achieve scalable performance on parallel machines. Efficient parallel implementation of such adaptive application is therefore a challenging task. In this paper, we compare the performance of and the programming effort required for two major classes of adaptive applications under three leading parallel programming models on an SGI Origin 2000 system ...

16 A Performance Evaluation of the Convex SPP-1000 Scalable Shared Memory Parallel Computer

Thomas Sterling, Daniel Savarese, Peter MacNeice, Kevin Olson, Clark Mobarry, Bruce Fryxell, Phillip Merkey

December 1995 **Proceedings of the 1995 ACM/IEEE conference on Supercomputing (CDROM) - Volume 00 Supercomputing '95**

Publisher: ACM Press, IEEE Computer Society

Full text available:  [pdf\(457.11 KB\)](#)
 [html\(2.67 KB\)](#) Additional Information: [full citation](#), [references](#), [citations](#)
 [ps\(780.23 KB\)](#)

17 Loop optimization for a class of memory-constrained computations

D. Cociorva, J. W. Wilkins, C. Lam, G. Baumgartner, J. Ramanujam, P. Sadayappan

June 2001 **Proceedings of the 15th international conference on Supercomputing**

Publisher: ACM Press

Full text available:  [pdf\(160.59 KB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

Compute-intensive multi-dimensional summations that involve products of several arrays arise in the modeling of electronic structure of materials. Sometimes several alternative formulations of a computation, representing different space-time trade-offs, are possible. By computing and storing some intermediate arrays, reduction of the number of arithmetic operations is possible, but the size of intermediate temporary arrays may be prohibitively large. Loop fusion can be applied to reduce memory ...

18 Informing memory operations: memory performance feedback mechanisms and their applications

Mark Horowitz, Margaret Martonosi, Todd C. Mowry, Michael D. Smith

May 1998 **ACM Transactions on Computer Systems (TOCS)**, Volume 16 Issue 2

Publisher: ACM Press

Full text available:  pdf(344.74 KB) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#), [review](#)

Memory latency is an important bottleneck in system performance that cannot be adequately solved by hardware alone. Several promising software techniques have been shown to address this problem successfully in specific situations. However, the generality of these software approaches has been limited because current architectures do not provide a fine-grained, low-overhead mechanism for observing and reacting to memory behavior directly. To fill this need, this article proposes a new class ...

Keywords: cache miss notification, memory latency, processor architecture

19 On the partitionability of hierarchical radiosity



Robert Garmann

October 1999 **Proceedings of the 1999 IEEE symposium on Parallel visualization and graphics**

Publisher: ACM Press

Full text available:  pdf(281.29 KB) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

The Hierarchical Radiosity Algorithm (HRA) is one of the most efficient sequential algorithms for physically based rendering. Unfortunately, it is hard to implement in parallel. There exist fairly efficient shared-memory implementations but things get worst in a distributed memory (DM) environment. In this paper we examine the structure of the IIRA in a graph partitioning setting. Various measurements performed on the task access graph of the HRA indicate the existence of s ...

20 GPGPU: general purpose computation on graphics hardware



David Luebke, Mark Harris, Jens Krüger, Tim Purcell, Naga Govindaraju, Ian Buck, Cliff Woolley, Aaron Lefohn

August 2004 **Proceedings of the conference on SIGGRAPH 2004 course notes GRAPH '04**

Publisher: ACM Press

Full text available:  pdf(63.03 MB) Additional Information: [full citation](#), [abstract](#)

The graphics processor (GPU) on today's commodity video cards has evolved into an extremely powerful and flexible processor. The latest graphics architectures provide tremendous memory bandwidth and computational horsepower, with fully programmable vertex and pixel processing units that support vector operations up to full IEEE floating point precision. High level languages have emerged for graphics hardware, making this computational power accessible. Architecturally, GPUs are highly parallel s ...

Results 1 - 20 of 60

Result page: [1](#) [2](#) [3](#) [4](#) [next](#)

The ACM Portal is published by the Association for Computing Machinery. Copyright © 2006 ACM, Inc.

[Terms of Usage](#) [Privacy Policy](#) [Code of Ethics](#) [Contact Us](#)

Useful downloads:  [Adobe Acrobat](#)  [QuickTime](#)  [Windows Media Player](#)  [Real Player](#)